



NCCS STORAGE POLICY

Version 1.1

Date 11/20/2009

The following policies apply to all projects with allocations on the National Center for Computational Sciences (NCCS) systems. The policies apply equally to all projects regardless of the machine on which they have time.

STORAGE TYPES

We have three storage types available to NCCS users: local storage, global storage, and archival storage. Each storage type has different purposes, storage limits, and purge/retention policies as described below.

Local Storage

A large, fast disk area intended for parallel access to temporary storage in the form of scratch directories may be provided on each system. This area is local to the specific system. This directory is, for example, intended to hold output generated by a user's job. Because of the scratch nature of the file system, it is not backed up and files are automatically purged on a regular basis. This storage type follows the Scratch Purge Policy. Files should not be retained in this file system and should be migrated to archival storage as soon as the files are not actively being used.

If a file system runs out of or gets low on space, user data on that file system will be subject to involuntary deletion without prior notice.

Quotas may be instituted on a machine-by-machine basis if deemed necessary.

Global Storage

There are three basic categories in global space: home space, scratch space, and project space. User areas in global space include home directories and individual directories in the center-wide global high performance file system.

Home Space

Home directories for each user are mounted on NCCS systems. This is online disk space intended for long-term, frequently used storage, and is backed up on a daily basis. This file system does not generally provide the input/output (I/O) performance required by most jobs. A user home directory has a 5 gigabyte quota.

Scratch Space

Individual user directories in the center-wide high capacity file system reside on large, fast disk areas intended for global (parallel) access to temporary storage in the form of scratch directories. These scratch directories are provided on each system. Because of the scratch nature of the file system, it is not backed up and files are automatically purged on a regular basis. Files should not be retained in this file system and should be migrated to archival space as soon as the files are not actively being used.

If a file system runs out of or gets low on space, then user data on that file system will be subject to involuntary deletion without prior notice.



Project Space

Project areas in global space include project directories for each project in the standard file system and in the center-wide high capacity file system.

Standard

Standard project directories are mounted on NCCS systems. This is online disk space intended for long-term, frequently used storage and is backed up on a daily basis. This file system does not generally provide the input/output (I/O) performance required by most jobs, but this area is backed up. The standard project directory has a 50 gigabyte quota.

High-capacity

The center-wide high capacity project directories are mounted on NCCS systems. This is online disk space intended for long-term, frequently used storage, but this area is **not** backed up. An individual project directory has a 1 terabyte quota.

Project space on the center-wide high capacity directories may be denied or revoked at anytime.

Archival Storage

High Performance Storage System, HPSS, is the archival storage system at the NCCS. Space on HPSS is intended for files that are not immediately needed. In the archival system, users have access to two types of allocation: user and project. User allocations belong to a single user. Project areas are shared between all users of the project.

User Space

An individual user has a quota of 2 terabytes aggregate space and 200 files. User areas in HPSS are created in /home.

Project Space

Project quotas have 45 terabytes aggregate space and 4,500 files. Project areas in HPSS are created in /proj.

Larger quotas may be granted for a project on a case-by-case basis; the projects must include storage needs in their proposals along with strategies for data migration after the project has ended. Even though HPSS is a very large storage system, space is not unlimited. Users must not store files unrelated to their NCCS projects on HPSS. They must also periodically review their files and remove unneeded ones. Quotas will be enforced.

Projects should not duplicate executables or data sets in each user's area but instead should set privileges to share a master copy. Contact the User Assistance Center for help in setting privileges to facilitate sharing.

Project data stored on HPSS will only be retained up to three months past the end of a project allocation. User data stored on HPSS will only be retained up to three months past user account termination. There is no lifetime retention. Projects are expected to migrate their data off HPSS when the project ends or before.



PURGE/RETENTION POLICIES

There are three purge/retention policy types: scratch policy, project policy, and home policy. Please note there is no lifetime retention for any data on NCCS machines. Projects and users are expected to migrate their data off of NCCS systems when projects and users are finished with NCCS systems. When a project ends, active users in that project will have 1 month to remove all their data off NCCS systems.

Scratch Policy

Scratch purge policies are intended to maintain file systems so there is always a large amount of scratch space available for executing jobs. Files are automatically purged on a regular basis. Any file older than 14 days is subject to deletion. If a file system runs out of or gets low on space, user data on that file system will be subject to involuntary deletion without prior notice.

Project Policy

Project purge policies are intended to reclaim space after a project ends. After a project ends, project data will be purged from all NCCS systems. If a project needs additional time to remove project data from NCCS systems, they must request an extension. When a project ends, users will have 1 month to remove all data from NCCS systems.

In cases where a project missed the deadline for an extension, NCCS reserves the right to delete all files and directories designated as project data.

Home Policy

Home purge policies are intended to reclaim space after a user account is disabled for any reason. When a user account is disabled, NCCS reserves the right to delete all files and directories in the designated home areas on archival, global, and local space.

In cases where project data is stored in home areas, a project representative must submit a request for data from the user's home area within 1 month of the user account being disabled. In cases where the deadline for an exception was missed, NCCS reserves the right to delete all files and directories designated as home data.

CONSEQUENCES OF ABUSE

Storage usage will be monitored continually. When time permits, offenders will be warned to clean up their space. Ignoring these warnings will result in loss of access privileges.



SUMMARY

Area	Actual Location	Policy	Quota	Purge Policy	BackUp?
Home Directory	/ccs/home/[username]	Home	5GB	1 month	Yes
NFS Project	/ccs/proj/[projectid]	Project	50GB	1 month post-project	Yes
Lustre Project	/tmp/proj/[projectid]	Project	1000GB	1 month post-project	No
Primary Scratch	/tmp/work/[username]	Scratch	-	14 days	No
Local Scratch	varies	Scratch	-	14 days	No
HPSS Home	/home/[username]	Home	2TB/200 Files	3 months	
HPSS Project	/proj/[projectid]	Project	45TB/4500 Files	3 months post-project	